

Requested Patent: GB1435927A

Title: DATA PROCESSORS ;

Abstracted Patent: GB1435927 ;

Publication Date: 1976-05-19 ;

Inventor(s): ;

Applicant(s): IBM ;

Application Number: GB19750022928 19730705 ;

Priority Number(s): GB19750022928 19730705; US19720272698 19720717 ;

IPC Classification: G06F7/38 ;

Equivalents: ;

**ABSTRACT:**

## PATENT SPECIFICATION

(11) 1435927

1435927 (21) Application No. 22928/75 (22) Filed 5 July 1973  
 (62) Divided out of No. 1435926  
 (31) Convention Application No. 272698 (32) Filed 17 July 1972 in  
 (33) United States of America (US)  
 (44) Complete Specification published 19 May 1976  
 (51) INT CL: G06F 7/38  
 (52) Index at acceptance

G4A 15A2 16D 16E4 2F6 2HX 4R 4X 5B AU  
 (72) Inventor MILTON JAY KIMMEL



## (54) DATA PROCESSORS

(71) We, INTERNATIONAL BUSINESS MACHINES CORPORATION, a Corporation organized and existing under the laws of the State of New York in the United States of America, of Armonk, New York 10504, United States of America, do hereby declare the invention, for which we pray that a patent may be granted to us, and the method by which it is to be performed, to be particularly described in and by the following statement:

The present invention relates to data processors.

Present-day commercial machines for recognizing printed characters and other patterns are almost invariably implemented in the form of hard-wired logic. General-purpose computers have been employed for many years for experimental and low-volume recognition, such as the design and simulation of character-recognition algorithms and the analysis of spacecraft photographs. The implementation of recognition functions in software form has many advantages: algorithms may be modified or replaced easily, the same hardware may be used for several different functions, to minimize idle time and the hardware and software are separately diagnosable, to name only a few. These advantages of course also apply to areas other than pattern recognition, and there is presently a trend toward the implementation of many different types of previously hard-wired functions in small, programmed data processors.

The major obstacle to the use of software recognition techniques in a commercial environment is speed. Since the circuits in a hard-wired machine can operate in parallel, presently available OCR (optical character recognition) and MICR (magnetic-ink character recognition) machines are capable of processing a thousand or more characters per second. Execution of the same algorithms on a conventional general-purpose computer, however, would lower this rate to the range of ten to one hundred per second. Careful and ingenious programming will increase this rate to some extent, but at the expense of greatly increased design time and at the limitation to

a small number of efficiently codable recognition techniques. The addition of special-purpose recognition micro-instructions for a general-purpose computer has also been proposed. While this achieves speeds generally comparable to those of hardware oriented recognition systems, it requires micro-instruction storage space which is always at a premium, it must be made less efficient in order to conform to architectural standards defined for other purposes, and it is limited to the implementation of a specific type of recognition algorithm.

What is required then is a "general-purpose" special-purpose recognition processor; a data processor which is special-purpose in that it is optimized to perform recognition and allied functions efficiently, yet which is general-purpose in that it is capable of implementing a wide range of recognition algorithms plus a number of other tasks, such as machine control, data management, and so forth.

According to the present invention there is provided a data processor including apparatus for performing a plurality of different voting-logic operations such apparatus comprising:

a plurality of cascaded registers;

data storage means for holding and accessing a plurality of data operands;

an instruction unit for accessing and interpreting program instructions, said instructions including a plurality of instructions respectively containing specifications of a corresponding plurality of said operands, and said program instructions including a further instruction containing a specification of a selectable one of said registers;

means coupled to said instruction unit and to said data storage means for entering successive ones of said operands into a first of said cascaded registers and for transferring the contents of the register sequentially from register to register so that for any register its new contents are a logical combination, with respect to a particular fixed value, of its former contents and the former contents of

50

55

60

65

70

75

80

85

90

95

the preceding register (or in the case of the first register, the new operand); and  
5 gating means coupled to said instruction unit and to all of the said cascaded registers for transferring to a common output line the contents of said selectable one register, so as to gate out one of a plurality of different voting-logic functions of said successive operands, depending upon which of said cascaded registers is specified by said further instructions.

10 The present invention will be described further by way of example with reference to an embodiment thereof as illustrated in the accompanying drawings, in which:

15 Figure 1 is a simplified block diagram of one type of system according to the present invention, designed in this case for pattern recognition;

20 Figure 2 is a data flow diagram of processor 200 of Figure 1;

Figure 3 shows the index unit 300 of Figure 2;

25 Figure 4 illustrates the logic unit 400 of Figure 2;

Figure 5 shows the details of the result-register unit 500 of Figure 2;

30 Figure 6 depicts the status registers 600 of Figure 2;

35 Figure 7 (shown as Figures 7A to 7J) describes a complete set of control logics for the processor of Figure 2.

#### Detailed Description of the Embodiment

35 The following description proceeds from an outline of a typical recognition system using a programmed processor to a description of the entire data flow of the processor. Certain subunits of the processor are then explained in greater detail. The operations 40 performed by each instruction in the processor's instruction set are then described, with particular emphasis on those instructions relevant to the invention. Finally, programming examples illustrate the use of the inventive concepts in representative applications.

#### Overall System

50 Figure 1 shows a representative character recognition system 100. A conventional mechanical transport 110 moves a character-bearing input document 120 past an optical scanner 130. Scanner 130 may provide, for instance, a vertical raster scan sweeping over one line of characters on document 120. Video detector 140 receives reflected light 55 from document 120 and converts it into an electrical signal representative of the light intensity reflected from document areas swept by the scanning beam. The signal on line 141 is conveniently digitized into black/white binary levels. It is also digitized in time, so that each binary signal or bit represents a predetermined-size area or cell on document 120. Apparatus for carrying out each of the above

functions is well known in the art. Other types of input devices, such as parallel optical 65 detector arrays, magnetic scanners, acoustic transducers or instrumentation outputs, may also be used in the present system.

Video processing unit 150 applies certain conventional preliminary transformations to 70 the signal on line 141. If system 100 is designed to recognize machine-printed or hand-printed characters, such functions may include electronic storage of character images, 75 character size normalization and position registration, and segmentation of adjacent characters. Circuitry for performing these functions is well known in the art, and may be of any suitable form. Output line 151 carries the transformed or "preprocessed" signal to processor 200 for recognition of its identity. Input/output adapter 160 receives an identification code signal from processor 200, and transmits it to a computer (not shown) or to any other suitable output device.

80 Although the primary task of unit 200 is to process input signals for recognition purposes, it also performs several other functions in system 100. Machine-control circuits 170, for instance, may receive condition signals from transport 110, relay them to processor 200, receive control signals back from processor 200, and transmit them to actuators for controlling the operation of the transport. Processor 200 may also receive control signals 85 from an external device through I/O adapter 160. One conventional use of such external control signals is to provide further control of transport 110, such as selecting a particular stacker for document 120. Such signals may also be passed through machine control circuits 170 to optical scanner 130, for formatting purposes. That is, it is frequently desirable to allow an external device to select 90 only certain parts of document 120 for scanning.

95 Back-up storage unit 180 holds several types of programs for recognition, machine control and diagnostics which cannot be simultaneously stored in the internal memories of the processor. The provision of storage 180 allows system 100 to operate off-line from a host central processing unit (CPU), or to make minimal demands upon the host CPU memory when it is operated on-line. Unit 180 100 may be a small disk drive or other form of conventional bulk memory.

#### Processor Configuration

110 Figure 2 is a data-flow diagram of processor 200. Before proceeding with a detailed description of Fig. 2, it may be helpful to explain certain conventions used therein. Each block in Fig. 2 contains either the full or an abbreviated name of a hardware unit represented by the block. This name, such as 115 "BAR", "PLR" and "Program Storage Module (PSM)", appears at either the left

side or the top of the associated block. Most blocks also contain a number at either the right side or the bottom of the block, representing the size of that block. The number "16" in the block labelled "BAR", for instance, indicates that the block is 16 bits wide. The number "64K x 16" in the program storage module block indicates that the capacity of this block is 64K (actually 65,536) words of 16 bits each. The lines between the blocks in Fig. 2 represent unidirectional data paths, each of which may carry multiple bits in parallel. The direction of data flow is shown by the arrows. Each block labelled "GATE" in Fig. 2 is capable of switching one or more data lines into the block below it or to its right. Such gates are common and well known in digital data processing equipment. To avoid undue complication in the description, the wiring for controlling the individual gates has been omitted from Fig. 2. Details of the gate-control functions will, however, be described hereinbelow.

Processor 200 is organized in two major subsections, an instruction unit 210 and an execution unit 250. Instruction unit 210 contains facilities for storing, interpreting and indexing program instructions; it has a relatively large amount of storage and wide data paths. Execution unit 250 includes components for receiving, storing and outputting data operands, and for executing logic functions. It contains smaller, faster storage facilities, and has narrower data paths. Timing unit 201 generates system timing signals and provides conventional synchronization for the various units of processor 200, over lines 202.

Instruction unit 210 is built around a program storage module (PSM) 220, an instruction decoder 230, a bank of status registers 600 and an index module 240.

PSM 220 is an on-line storage module capable of containing a set of instructions for at least one program. It is here shown as having 65,536 words of sixteen bits each. With respect to the program instructions, PSM 220 appears to be a read-only memory; that is, no instruction in a program can directly modify the PSM contents. Program load register (PLR) 221, however, provides facilities for loading different programs from back-up storage 180, over line 181. PLR 221 also provides signals to gate 222 for initializing program address register (PAR) 223 while the program is being loaded. Control signals for program loading originate in an external device (not shown).

During execution of a program stored in PSM 220, addresses in PAR 223 may be derived from one of three sources, depending upon the setting of gate 222 by instruction decoder 230. For the execution of instructions in sequential locations of PSM 220, incrementor 224 adds a constant value (i.e., one) to the output of PAR 223 and returns it to one

of the selectable inputs of gate 222. In order to modify the normal sequence of execution, branch address register (BAR) 225 may also be gated to PAR 223 through gate 222. It should be noted at this point that BAR 225 is capable of holding the single branch address throughout the execution of an indefinite number of other instructions. The importance of this fact will be explained at a later point. BAR gate 226 selects various portions of the branch address either from the program output register (POR) 228, or from fetch output register (FOR) 252 of execution unit 250, depending upon the type of instruction being executed. Thus, the branch address in BAR 225 may be specified in whole or in part by an instruction from PSM 220 and by the result of a computation in execution unit 250. The third source of addresses for PAR 223 is branch-and-link register (BLR) 227, which receives an output from incrementor 224 and which may be selectively transferred through gate 222 into PAR 223, to provide return linkages for subroutines.

POR 228 holds the current instruction which has been addressed by PAR 223. The eight most significant bits of the sixteen bits in POR 228 are transferred to operation code register (OPR) 229, although some of these bits are not decoded as operation codes for certain instructions. The least significant ten bits provide one of the inputs to index gate 241. Either eight or twelve bits of POR 228 are made available to BAR gate 226 for entry into selectable portions of BAR 225.

Instruction decoder 230 contains conventional forms of combinatorial and sequential logic for interpreting the operation codes stored in OPR 229. Output lines 203 then carry time control signals to the various gates of processor 200 for routing data through the appropriate paths at various times within an instruction execution cycle. Decoder 230 also produces output signals on lines 232 to status registers 600. Registers 600 also receive inputs over lines 204, from various units in processor 200, and produce output signals on lines 205, as will be described in connection with Fig. 6. Detector 231 detects all-zero operand addresses in POR 228, since these addresses modify the effect of certain instructions.

Index unit (IU) 300 of module 240 provides operand-address modification for indexed instructions. It also provides an arithmetic capability required by certain functions of the processor. Briefly, gates 241 and 242 each provide one operand to IU 300, which forms a variable-modulus sum in summation register (SR) 243. Gate 241 receives ten-bit operands from POR 228 and eight-bit operands from result register two (RR2) 520 of execution unit 250. Gate 242 receives inputs from index register (XR) 244, and from FOR 252 of the execution unit.

70

75

80

85

90

95

100

105

110

115

120

125

130

Execution unit 250 is organized around a data storage module (DSM) 260 and an associated stack storage module (SSM) 270, a set of input and output, or "external" registers (EXR) 280, a logic unit (LU) 400, and a set of result registers (RR) 500. These components are employed to receive, store and output system data, and to perform logic operations upon these data. 70

5 Data from data input register 253 may be gated into DSM 260 by control signals transmitted to gate 261. Data address register (DAR) 262 controls the DSM location at which these data are stored; since DSM 260 has 2,048 8-bit locations, DAR 262 is 11 bits wide. The address in DAR 262 may be derived from several different sources, depending upon the setting of gate 263. The most significant bit ("H" bit) of this address is 75

10 always derived from a line 205 from one of the status registers 600, as will be explained. Most instructions specify that the low-order 10 bits of the address are derived from SR 243, over line 206. These addresses are ultimately contained in instructions in POR 228, and are transferred to SR 243 either directly or in modified form by index module 240. Input data to processor 200 from external sources, however, are normally addressed as a 80

15 queue in contiguous locations of an area of DSM 260. For this purpose, top register (TOP) 264 contains an address which is adjacent the address of that location which had last received a data byte. TOP 264 is updated by incrementor 265 via gate 266. Those instructions which have the effect of writing video data into DSM 260 (in a manner which will be detailed at a later point) also control gate 266 to place in TOP 90

20 264 the address which is next higher to that contained in DAR 262. TOP 264 may also be initialized from SR 243, by switching line 206 through gate 266. In a like manner, bottom register (BOT) 267 contains the address of the location adjacent the last data byte which had been read out of DSM 260. BOT 268 is updated by incrementor 265 during the execution of an instruction which 95

25 reads video data from DSM 260. BOT 267 may also be initialized from SR 243 by switching line 206 through gate 268. Data are read out of DSM 260 into data output register (DOR) 269. Indexed instructions require that a string of zeros be supplied to DOR 269, regardless of the contents of the actual DSM location addressed by DAR 262, when certain conditions are fulfilled in the computation of a modified address by index unit 300. This effect is achieved by disabling gate 26A and resetting DOR 269 for those instructions. 100

30 The storage and use of intermediate computational results is facilitated by a stack storage module (SSM) 270 which is essentially parallel with DSM 260. Data are 105

35 pushed into SSM 270 from DIR 253 by opening gate 271 while simultaneously closing gate 261, and thereafter incrementing stack address register (SAR) 272. Data are popped from SSM 270 through gate 291 and SAR 272 is then decremented. SSM 270 is capable of storing 16 8-bit bytes. 110

40 Communication from DSM 260 to instruction unit 210 is mediated by fetch output register (FOR) 252 and lines 207. Gate 251 assembles two consecutive 8-bit bytes into the 10-bit format required by instruction unit 210. 115

45 Logic module 290 performs one of a number of selectable operations upon data from gates 291 and 292, and enters the result in result register one (RR1) 510. One input to logic unit (LU) 400 is selected by gate 291 from SR 243, DOR 269, SSM 270, or one of the external registers (EXR) 280. The other input, when used, is obtained by gate 292 from result register two (RR2) 520. A byte from RR1 510 may be gated to either RR2 520 or to one of the EXR's 280. The selection of a particular EXR (such as EXR 32 282) for the receipt or transmission of a data byte is indicated by certain instruction bits from POR 228, and is performed in a conventional manner by selection logic associated with the EXR's. The contents of RR2 520 may be transmitted to IU 300 through gate 241, to LU 400 through gate 292 or to DIR 253 through gate 254 or in a modified form to result register three (RR3) 530. RR1 510 may also be gated directly to DIR 253 through gate 254. RR3 530 may be gated only to DIR 253, again through gate 254. 120

50 Gate 254 also controls the input of video data from line 151 into DIR 253, through an external register specialized for that purpose, the video register (VR) 281. The manner in which the loading of video data occurs during the execution of other instructions will be explained in connection with Fig. 6. 125

#### Index Unit

Index unit 300, Fig. 3, receives signals from gate 241 on lines 301 and signals from gate 242 on lines 302. When a signal on line 203 indicates that the operand address of an instruction is to be indexed, IU 300 produces a cyclic sum output on line 304. The modulus of the sum is variable, and is controlled by signals on lines 305, 306 and 307. Line 308 produces a carry signal indicating that the modulus selected by lines 305-307 has been exceeded during the summation operation. 115

The reference numerals 310, 320, 330, 340 and 350 respectively represents the first, second, eighth, ninth and tenth stages of a 10-bit adder network. The third through seventh stages of this network are identical 120

60 65

to stages 310 and 320. A logic "zero" signal on line 203 indicates that the signal on line 301 is to be moved directly to output line 304, without indexing. Accordingly, a "zero" level on line 203 enables AND gates 311, 321, 331, 341 and 351 through inverters 312, 322, 332, 342 and 352, to pass the signals on lines 301 directly through OR gates 313, 323, 333, 343 and 353 to the corresponding output lines 304.

When line 203 carries a logic "one", AND gates 314, 324 and 334 transmit the sum outputs 315, 325 and 335 of conventional full adders 316, 326 and 336 through OR gates 313, 323 and 333 to output lines 304. Conventional interstage carry lines 317 and 327 are connected between the full adders, as in conventional practice.

A logic "one" signal on line 307 then indicates that the modulus of the sum signal on lines 304 is  $2^{10}$ , i.e., 1,024. (When any one of the lines 305-307 carries a "one", the other two are always at the "zero" level). Under these conditions, the zero level on line 305 enables AND 344 through inverter 347 to transmit the sum output 345 of adder 346 through OR 343 to output lines 304. Line 307 enables AND 354 to pass the sum output 355 of adder 356 through OR 353 to the high-order position of lines 304. Line 307 also enables AND 361 of carry logic 360 to pass the carry output 357 of full adder 356 through OR gate 362 to carry output 308. That is, when lines 305-307 call for a 10-bit sum, index unit 300 operates as a conventional 10-bit adder with an indication of a carry out of the tenth position.

When, however, a "one" appears on line 306 (and "zeroes" on line 305 and 307), AND gate 358 is enabled through inverter 359, while AND gate 354 is disabled. Since AND 358 is connected to one of the lines 301 instead of to the output 355 of full adder 356, the most significant bit of input 301 is passed to the most significant output bit 304 directly rather than being added to the most significant bit of input 302. Line 306 also conditions AND gate 363 to transmit the carry signal on line 347 from adder 346 through OR 362 to carry output 308.

With a "100" pattern on lines 305-307, AND gate 358 remains enabled, but OR 343 receives a signal from AND 348, rather than AND 344. Line 305 also conditions AND 364 to produce the carry signal on line 308 from the carry output 337 of full adder 336. Under this condition, then, the low-order eight bits of line 301 are added to the low-order eight bits of lines 302, while the two high-order bits of line 301 are transmitted directly to the two-high order bits of line 304.

It may now be seen that index unit 300 operates as a variable modulus adder which

also has a variable base. That is, it is a modulo-1024 adder when line 307 is active; when line 306 is active, the modulus is 512, and the possible range of sums is either 0-511 or 512-1023, depending upon the high order bit of line 301. A carry signal is produced on line 308 when the sum generated by stages 310-340 exceeds 511, regardless of the value of the high-order bit of line 301. When line 305 is active, the addition modulus is 256, and the range of possible output numbers is 0-255, 256-511, 512-767 or 768-1023, depending upon the values of the two high-order bits of lines 301. Again, line 308 has a carry signal when the sum generated by stages 310-330 exceeds 255, regardless of the values of the two high-order bits of line 301.

It will now be understood how index unit 300 provides a variable-modulus, variable-base wrap-around indexing function for data operand addresses. More or fewer stages may of course be provided in obvious ways, and other methods of implementing the above function will appear to those skilled in the art. The signal sources for lines 305-307, and the disposition of the carry signal on line 308, will be discussed in connection with Fig. 6.

#### Logic Unit

Logic unit 420, Fig. 4, accepts signals on lines 402 from gate 241 (Fig. 2) and on lines 402 from gate 242, and provides output signals on lines 404 in accordance with one of a number of operations specified on lines 203 from instruction decoder 230. The signals on lines 203 and 404 also cooperate to produce a condition signal on line 403.

As in common practice, instruction encoder 410 provides signals on line 411 indicating which logic operation is to be performed for an instruction decoded by instruction decoder 230. These signals control conventional logic circuits 420. A signal on line 412 additionally conditions exclusive-or (XOR) gates 421 in order to transmit either true or complement signals from line 401 to line 422. Logic circuits 420 comprise eight controllable circuits, one for each of the bits on lines 402 and 422, and are capable of performing the following operations: An "OR", "AND" and "Exclusive-OR" upon corresponding bits of each of the lines 402 and 422, and direct transfer of the bits from lines 422 to lines 404. Other conventional logic and/or arithmetic functions may of course also be provided. Condition detectors 430 are coupled to output line 404 for detecting various bit combinations resulting from the logic operations performed by circuits 420. Detector output lines 431-434 may, for instance, provide signals respectively indicating that all the output bits are all "zero", all "one", not all "one" and

70

75

80

85

90

95

100

105

110

115

120

125

not all "zero". These lines are coupled to the inputs of AND's 435-438 which also receive signals from instruction encoder 410 via lines 414-417. Inputs 414-417 select one of the conditions which may have been detected by detector 430 for transmission through OR gate 438 to condition-output line 403.

An instruction signal on line 203, for instance, which indicates that logic circuits 420 are to perform an "AND" operation will produce a signal on line 414, so that an all-zero result on line 404 will produce a condition signal on line 403. Similarly, an "OR" type of instruction will produce a signal on line 415, so that line 403 will produce a signal when the output bits are all one. As is indicated in Fig. 4, other types of instructions may call for other bit configurations to produce an output on line 403. It is also possible, of course, that two different instructions from decoder 203 which produce identical signals on lines 411 may produce different signals on line 413. An "And Mask Not Zero" instruction, for example, performs a logic AND of two operands, but activate line 417 instead of line 414, so that output 403 produces a condition signal when the result is not all "zero". Although it has not been found necessary here, it would also be possible to provide a specific instruction for selecting among the various conditions to be tested.

The signal on line 403 controls the modification of the instruction-execution sequence, in conjunction with certain bits of status registers 600, which will be described in connection with Fig. 6.

#### Result Registers

The purpose of result register unit 500 Fig. 5, is to perform voting or counting operations upon successive input operands, and to gate the results to one of several output lines. Unit 500 also performs "shift-load" operations.

Each result register unit 510, 520 and 530 contains a bank of eight AND gates 511, 521 and 531, and an eight-stage trigger register 512, 522 and 532. Each trigger register is of the conventional type in which each stage may be individually set to a logic "one" level, but the stages are reset to a "zero" level in parallel by a single "clear" input 513, 523 and 533. Such registers also commonly employ a "clock" line (not shown) for synchronization purposes.

As will be explained hereinafter, certain instructions in the instruction set of processor 200 cause instruction decoder 230 (Fig. 2) to produce a "push" signal on one of its output lines 203. This signal conditions inputs 534 of AND gates 531 to set those stages of register 532 for which a logic "one" signal appears on both the inputs 535, which are connected

to the eight bit positions of the logic-unit output lines 404 and the outputs 526 of RR2 520. Outputs from the individual stages of register 532 are collected at outputs 536, and are routed to other parts of processor 200 on lines 537 as shown in Fig. 2.

AND gates 521 of result register 2 (RR2) 520 are conditioned by push inputs 524 and by outputs 516 of RR1 510, acting through shift gates 540. Gates 521 are then fully enabled by inputs 525, which also are connected to the eight bits of lines 404 to enter logic "ones" into particular stages of trigger register 522. The output 526 of each stage of register 522 is transmitted within the processor 200 as is shown in Fig. 2. Outputs 526 and push inputs 534 then proceed to AND's 531 of result register 3 (RR3) 530 as described.

Inputs 515 are also coupled to the respective bit positions of lines 404 and to push line 203 to enter data from LU 400 into trigger register 512. (Inputs 518 may be considered to be active for the moment). RR1 output lines 516 are coupled to line 517 and routed as shown in Fig. 2. Outputs 516 also proceed through shift gates 540 to RR2 520.

Additional stages may be coupled to RR3 530 in the same manner in which it is connected to RR2 520, if desired.

The operation of result-register unit 500 for voting logic is as follows. Considering for the moment only a single bit position of each register, and disregarding shift gates 540, assume that all registers 512, 522, and 532 have been initially cleared to "zero" by signals on lines 513, 523 and 533. A push signal on lines 514, 524 and 534 will not change the contents of any trigger registers if the corresponding data bit on inputs 515, 525 and 535 is a "zero". If the latter inputs carry a "one", however, it will be entered only in register 512, and not in registers 522 and 532. A "zero" at inputs 515, 525 and 535 on a succeeding machine cycle will have no effect on the register contents, since register 512 can only be reset by input 513, and ANDS 521 and 531 are not fully enabled. A similar result obtains when push inputs 514, 524 and 534 are zero, whatever the state of the data inputs. When both the push inputs and the data inputs are logic "one", however, AND 521 is fully enabled by inputs 524, 525 and 516 (via output 546) to enter a "one" into register 522. Because of a delay inherent in the operation of the trigger registers, output 526 is not yet capable of providing an input to AND 531. Thus, at the end of two instruction cycles, registers 512 and 522 both contain a "one", while register 532 contains a "zero". On a third instruction cycle, the presence of either a zero on push inputs 534, or a one on 534 and a zero on data input 535, will leave the contents of all registers unchanged. A logic "one" on the inputs 534 and 535, however, will allow the "one" on line 526 to be

65

70

75

80

90

95

100

105

110

115

120

125

5 entered through AND 531 into register 532. Any additional stages coupled to register 530 would operate in a similar manner. That is, when the "push" line is active, the "one" bits of successive data operands are propagated to successive stages of register unit 500.

10 Gate 254 (Fig. 2) may then select one of the lines 517, 527 and 537 under the control of a signal from instruction decoder 230. A logic 15 "one of n" function of two or more successive operands, for instance, may be entered into DIR 253 merely by selecting lines 517 to be transmitted through gate 254. The logic functions "two of n" and "three of n" may be selected by respectively gating either lines 527 or 537 into DIR 253 after the passage of n successive operands.

20 Output lines 517 and 527 may also be gated to other points within the processor 200.

25 Shift gates 540 and input lines 518 cooperate to perform another function, known as "shift-load". The significance and use of this operation will be described in connection with the instruction set. During the execution of voting-logic instructions, shift gates 540 pass the signals on lines 516 directly to corresponding output lines 546 without any modifications.

30 When a shift signal from decoder 203 is present on input 544, gates 540 selectively shift the signals on lines 516 by one, two or four places to the left or right, depending upon the state of the shift signal. Circuits for performing this function are well known. The shift is of the "end-around" type, wherein any bits shifted off the left end of gates 540 are re-entered on the right end, and vice versa. The bits on line 527 are then passed through LU 400 and selectively re-entered into register 512, after registers 512 and 532 have been cleared. The selection is accomplished by energizing only certain ones of the inputs 518, in accordance with particular bits in the current instruction. The instruction then places selected bits from another operand on line 404 into other positions of register 512, again under the control of inputs 518. If RR1 originally contains an operand J having bits (J0 J1 J2 J3 J4 J5 J6 J7), and an instruction calls for the shift-left load-right of an

40 operand K having bits (K0 K1 K2 K3 K4 K5 K6 K7) by two places, the final result register contents would be:

45 RR1: J0 J1 J2 J3 J4 J5 K6 K7  
 RR2: J2 J3 J4 J5 J6 J7 J0 J1  
 RR3: 0 0 0 0 0 0 0 0

50

#### Status Registers

55 Status registers 600 are groups of single-bit latches whose function may be broadly stated as that of retaining specific items of information over the execution of a number of instructions within processor 200. They may also be conceptualized as extensions to OPR

229, to provide additional bits in the operation codes of certain instructions which alter the operations performed by those instructions. Such alteration greatly multiplies the apparent size of the instruction set, yet does not significantly increase the amount of storage required for operation codes. As a general statement, the status registers 600 may be said to embody the heart of the present invention.

65 Registers 600 receive inputs 232 from instruction decoder 230 and inputs 204 from other units within processor 200. The particular instructions which produce signals on lines 232 will be discussed in greater detail in connection with Fig. 7. Output lines 205 control gating and other functions within processor 200.

70 The "discard result register" register 610 indicates whether the contents of RR1 510 have been stored elsewhere, and may thus be overwritten by subsequent instructions. A group of instructions, referred to as "store RR" instructions, set latch D, which disables AND gate 611 via inverter 612, since any instruction which stores RR1 510 makes its contents available for later instructions. As has been explained, however, any "logic" instruction executed by LU 400 produces an output on lines 404 which modifies the contents of RR1 510; therefore, these instructions reset latch D through OR 613, which enables AND 611 through inverter 612 so that a subsequent "load" instruction will transmit a signal on an output line 205 which causes gates 254 and 271 to transfer the contents of RR1 510 through DIR 253 to SSM 270. A "load"-group instruction by itself is prevented from activating AND 611 by delay 614, although a second "load" instruction will cause a transfer. That is, "store RR", "logic" and "load" instructions affect the state of latch D, but only "load" instructions use the state of latch D to transfer RR1 510 to SSM 270.

75 85 90 95 100 105 110 115 120 125 The "high-bit" register 620 specifies the value of the most significant bit of DAR 262, via gate 263. The occurrence of a "set addressing status" (SAS) instruction from decoder 230 enables AND gates 621 and 622 to receive a control bit from the SAS instruction located in SR 243. A "one" bit will set latch H through AND 621, while a zero bit will reset latch H through inverter 623 and AND 622. An output line 205 then transfers the value of this bit through gate 263 to DAR 262, as shown in Fig. 2. Thus, register 620 specifies whether subsequent instructions will address the top half or the bottom half of DSM 260.

"Data-control" register 630 maintains a record of the amount of video data present in DSM 260, in order to avoid writing into the data queue when it is full or reading from it when it is empty. Register 630 also allows data to be read automatically from video processor 150 in parallel with the execution

of other instructions within processor 200. Whenever processing circuits 150 are ready to transmit a video byte from register 281 into DSM 260, a signal on line 151 sets latch V to condition AND 631. Two other conditions, however, are required before this request may be honored. The first condition is that the instruction currently being executed by execution unit 250 does not require an operand to be read out of or written into DSM 260. Such instructions are grouped together as "video service" instructions in decoder 230 and provide an output signal on a line 232 to provide another input to AND 631.

15 The other condition is that DSM 260 have an available memory location for the video data. This condition is transmitted from overflow latch Q to AND 631 through inverter 632. The output of AND 631 is transmitted to gates 253 and 254 to allow VR 281 to be transferred to DIR 253 and thence to DSM 260. The output of AND 631 is also fed back to AND 633. This gate then sets latch Q if comparator 634 detects that the contents of 20 TOP 264 is equal to the contents of BOT 267. That is, an equality in the contents of registers 264 and 267 occurring immediately after the loading of a video byte indicates that DSM 264 has wrapped around to the point 25 at which the video string begins, and can therefore not accept any further data. As will be explained in greater detail hereinafter, the "load video" (LV) instruction has the effect of destructively reading a video byte from DSM 260. This occurs when AND gate 635 passes a signal to gate 26A to read a byte from DSM 260 into DOR 269. Underflow latch U, however prevents this transfer via inverter 636 when no video data is present in DSM 260.

30 Underflow latch U is set through AND 637 and OR 638 by the output of AND 635 whenever comparator 634 indicates equality between the contents of registers TOP and BOT. This condition signifies that a previous 35 transfer has collapsed the data string to the point where the beginning and ending addresses are the same. Latch U is reset by any of the video-service instructions when latch V is set, since these instructions allow a video byte from VR 281 to be loaded into DSM 260, and V indicates that a byte is ready for loading. The output of AND 635 also resets overflow latch Q through OR 639, since a transfer out of DSM 260 creates at least one free location 50 for the writing of a new video byte. Latches U and Q may be initialized to indicate an empty data string by a special instruction ZTB acting through OR's 638 and 639.

55 Register 640 controls the execution of conditional branching in instruction unit 250. Its operation is initiated by a "condition branch" (where "condition" is a verb, not a noun) instruction having the mnemonic CB. The detection of a CB operation code in register 229, Fig. 2, causes instruction decoder 60 65 to store bits 4-15 of the CB instruction in BAR 225, and to emit a signal on one of the lines 232 to set limit latch L, Fig. 6. This line also sets mask latch M through OR gate 64. Thereafter, a subsequent instruction which produces a signal on condition line 403 (from condition detector 430, Fig. 4) causes AND 642 to set branch latch B, the output of which effects the branch operation by controlling gate 222, Fig. 2, to transfer the contents of BAR 225 to PAR 223. Latch B is then reset for subsequent instructions by a signal on a line 202 from timing circuits 201. The output of latch B also resets latches L and M through ORs 643 and 644, in order to prevent instructions subsequent to the branch execution from again transferring to the same branch address.

70 Latches L and M may be reset through ORs 643 and 644 even when the branch has not been executed. This action is desired when the instruction stream contains any one of a number of instructions which may be termed "de-condition branch" instructions. Such instructions produce a signal on one of the lines 232 from instruction decoder 230. The action of these instructions is to destroy any further access to the address placed in BAR 225 by any preceding CB instruction. Certain other instructions, known as "load" 75 instructions, prohibit the execution of a branch during the execution of one or more instructions, but still allow a branch address specified by a previous CB instruction to be executed when certain further conditions have been met. The "load" instructions reset latch M through OR 644. The simultaneous presence of three further conditions then allows subsequent instructions to execute a branch to the address specified by the previous CB instruction, by setting latch M through AND 645 and OR 641. One of the conditions required by AND 645 is a "set" state of latch L: that is, register 640 must not have been completely deconditioned through OR 643. Another condition is the occurrence of an all-zero address in SAR 272, which is sensed by zero detector 646. As will be explained, that is necessary in order to prevent a branch address which had been set up at one level of a logic function from being taken at a different level of the function. The third condition required by AND 645 is the existence of an instruction in the group named "stack logic" instructions. Those instructions in the instruction set of processor 200 which belong to the "de-condition", "load" and "stack logic" groups are identified in Table I, below.

80 90 95 100 105 110 115 120 125 130 "Carry" register 650 is responsive to an SXS (set indexing status) instruction to provide signals on lines 305-307 for controlling index unit 300, Fig. 3. The detection of an SXS instruction by instruction decoder 230 produces a signal on one of the lines 232 which enables AND gates 651-654 to accept

two of the four SXS control bits on lines 204 from SR 243, Fig. 2. These two bits will then set or reset latches C0 and C1 through AND's 651-654 and inverters 655 and 656. The four possible bit combinations in these latches are then decoded by decoder 657 to produce a signal on only one of the lines 305-307.

"White-fill" register 660 controls the effect of a wraparound which occurs in Index Unit 300. Register 650 receives control bits from an SXS instruction and carry signals from line 308 of IU 300. It produces a signal on a line 204 for controlling gate 26A of DOR 269, Fig. 2. Again, the detection of an SXS instruction by instruction decoder 230 produces a signal on a line 232 which conditions AND's 661-664 to receive two SXS control bits from SR 243. These bits set or reset latches W0 and W1 via AND's 661-664 and inverters 665 and 666.

When neither latch is set, OR 667 is prevented from producing an output under any condition, since neither AND 668 nor AND 669 is enabled. Under this condition, the data stored in DSM appears to repeat itself cyclically at the modulus specified by lines 305-307. The data may be stored in one of several different areas in DSM 260, however, because of the pass-through action of IU 300 and the high-bit logic 610. When line 307 is active, for instance, a 1024-byte data queue may be located either at DSM addresses 0-1023 or at 1024-2047, depending upon the state of latch H of register 610. A 512-byte data queue wraps around between DSM addresses 0-511, 512-1023, 1024-1535 or 1536-2047; its location depends upon both latch H and the high-order bit of the absolute operand address, since the latter bit passes through IU 300 unchanged when line 306 is active. Similarly, a 256-byte area may be placed in any one of eight blocks of DSM 260.

When latch W1 is set, however, AND 668 and OR 667 operate to disable gate 26A whenever a carry signal on line 308 indicates that an instruction has caused IU 300 to exceed the modulus selected by lines 305-307. When latch W0 is set, the absence of a carry from IU 300 opens gate 26A. When gate 26A has been thus opened, the operand appearing in DOR 269 contains all zeros, instead of the contents of the DSM location passed from SR 243 through gate 263 to DAR 262. The net effect of this action causes DSM 260 to appear to an indexed instruction as having an area whose size is controlled by lines 305-307, and whose location is specified by latch H and possibly by one or two bits of the operand address contained in the instruction; outside that area, all data operands appear to contain zeros. The fourth possible combination, wherein both of the latches W0 and W1 are set, is not used in the present embodiment.

#### Control Gating

As has been stated, instruction decoder 230 and status registers 600 provide control signals to the units of Fig. 2 to interconnect them along specific data paths during the execution of each instruction from PSM 220. Fig. 6 shows the detailed implementation of logic circuits for status registers 600. The construction of a detailed implementation of the remaining control logic is straightforward, and will be obvious from the following operational description of a complete instruction set for processor 200.

To this end, Fig. 7 contains an APL description of the gating logic required for each instruction in the set. The APL notation is explained in, for example, "APL/360 Primer", IBM Systems Reference Library Form No. H20-0689-0 (Sept. 1969). In the description, the origin will be taken to be 0; then, the function 8, e.g., would expand to "0 1 2 3 4 5 6 7".

Fig. 7 first lists several common gating functions which are shared among many instructions. The first requirement for any instruction, for instance, is to fetch it from PSM 220. The function FETCH, therefore, first moves to POR 228 the contents of a location in PSM 220 specified by the value (to the base "two") stored in PAR 223; then, OPR 229 receives bits 0-7 of POR 228. Note that the names of the data sources and destinations in Fig. 7 correspond to the abbreviations used in the drawings and in the preceding description.

The remainder of Fig. 7 lists each instruction in the set in the alphabetical order of its mnemonic. The "And" instruction A, for example, is first obtained with the FETCH function, explained above. Since the operand for this instruction is not indexed, the NOINDEX function causes gates 241 and 242 to place bits 6-15 of POR 228 into SR 243, and catenates the H bit and the contents of SR 243 into DAR 262. DFETCH then gates the contents of the DSM location specified by DAR 262 into DOR 269. Next, RR1 510 is gated to RR2 520 and the AND function of the RR2 and DOR bits is returned to RR1 510. Branching is enabled as a part of the execution of this instruction by setting the B latch if latch M is set and if the result of the AND is all zeros. RESET then zeros the contents of RR2 and RR3, and latch D is reset to zero. Finally, NEXT obtains the address of the next instruction, from BAR 225 if B had been set; otherwise, the next sequential instruction is selected from DSM 220.

Four instruction formats are employed. Those instructions which specify a branch address have an operation code in bits 0-3 and an address value in bits 4-15. These instructions are B, CB and BAL. Instructions which specify an operand to be obtained from DSM 220 have an operation code in bit

positions 0-5 and an address value in positions 6-15. Examples are the A, L, LR4, OX and XORS instructions.

5      Immediate-operand instructions, such as AMZ, CME and ZR, have an operation code in bits 0-7 and a byte of immediate data in bits 8-15. Control and miscellaneous instructions, such as READ, WRT and SXS also have an operation code in bits 0-7; bits 8-15 of these instructions contain up to eight control bits for various purposes.

Fig. 7 shows which of the instructions produce signals on the various lines labelled 232 in Fig. 6. The "Store RR" instructions which contain "D<-1" in their description. Likewise, the "logic" and "Load" instructions contain "D<-0". "Load" instructions also contain "M<-0", since they reset latch M, but not affect latch L. "Video Service" instructions may be identified by the VIDEO common function. The "De-condition" instructions all reset latches L and M that is, "L<-M<-0". "Stack Logic" instructions are those which set latch M if latch L is set and if the address in SAR 272 is zero. These instructions, therefore, contain the descriptive notation

"M<-M v L&0=v/SAR".

#### Operation

30     The operation of processor 200 will now be described in detail for the execution of those instructions which are critical to the present invention.

#### Branch Conditioning and Execution

35     The "Condition Branch" (CB) instruction is different in several respects from the "Branch on Condition" (BOC) or "Conditional Skip" instructions of conventional data processors. BOC instructions contain branch addresses 40 which become effective during the execution of the branch instruction itself, based upon the existence of some previously determined condition. The specification of the particular condition to be tested is contained in the BOC instruction. A branch is either taken at 45 a specific time in the program or it is never taken.

50     The present CB instruction, on the other hand, cannot of itself alter the sequence of program instructions. Rather, it operates to convert certain subsequent instructions into BOC instructions, and establishes an address 55 to which they may branch; that is, CB merely conditions subsequent instructions to operate as branch instructions if certain conditions are fulfilled during their execution. The conditions which actually effect a branch are contained in the subsequent instructions, and not in the CB instruction. The AMNZ and AMZ instructions described in Fig. 7, for instance, produce identical results when a branch has not been conditioned, or when it has been

masked or de-conditioned as explained in connection with Fig. 6. But, when a branch has been conditioned and not masked or de-conditioned, AMZ transfers control to the branch address when logic 420 produces all zeros, and AMNZ transfers control when RR2 520 does not contain all zeros.

65     The branch which is enabled or conditioned by the CB instruction cannot be executed by all subsequent instructions. Some instructions, such as SAS, cannot cause a branch under any condition resulting from its execution. Instructions such as LK and SLS implicitly operate as "de-condition branch" instructions, in that they completely remove the conditioning established by a previous CB. "Load" 70 instructions are those which implicitly remove the CB conditioning during their own execution and for subsequent instructions, until a "stack logic" reconditions the branch. The re-conditioning provided by stack-logic 75 instructions (AS or OSC, e.g.) is itself conditional, since it takes place only when SAR 272 80 contains all zeroes.

85     Referring to Figs. 2, 6 and 7, the CB instruction itself activates the following gates. Initially, FETCH allows decoder 230 to decode the operation code, as has been explained above. Next, latches L and M are both set. Bits 4-15 of the instruction are then transferred to BAR 225, through gate 226. VIDEO services a video request if present, since DSM 260 is not otherwise used during the CB instruction. SAR 272 is then set to "0000". Finally, NEXT sets up the next sequential DSM address for the next instruction, since latch B was not set by the CB instruction. At this point the branch address specified in 90 bits 4-15 of the CB instruction has been entered and retained in BAR 225, but no branch has been taken under any conditions 95 which may have been produced by previous instructions.

100    Table I shows a program segment which illustrates the effect of a CB instruction, which appears as statement #2. Statement #3 performs a logical AND of the symbolic operands MAF2 and PC. From the control-gating description in Fig. 7, the instruction at statement #3 sets latch B if the logic AND of these two operands is all zeros, since latch M had been set by the CB instruction. If statement #3 does set latch B, the branch address "#57" is transferred from BAR 225 to PAR 223, and the next instruction executed is the BAL at statement #57. If latch B was not set, then statement #4 is executed next. Then, if that instruction produces an all-zero result, statement #4 will set latch B, and control will be transferred to statement #57. Similarly, an all-zero result in any of the statements #5 through #14 will effect a branch to statement #57, without requiring any further branch instructions.

110    The above statements all performed a logic

65

70

75

80

85

90

95

100

105

110

115

120

125

AND function, and effected a branch whenever the result contained all zeros. Other types of instructions, however, cause the branch to be taken under other conditions.

5 Other instructions have yet other conditions; ADD, for instance, tests bit 15 of POR 228 and bit 1 of SR 243; TOX tests the high-order bit of DOR 269 and branches if it is a "one". These examples point out the facts that the conditions themselves may involve a separate computation, that the conditions to be tested may be obtained from different sources for different instructions, and that data or latch settings may be employed in addition to a direct computational result.

10 Proceeding through Table I, no branch can be taken at statement #15, since Fig. 7 shows that an LC instruction resets latch M. Instructions which reset latch M without resetting latch L are termed "load" instructions, since they perform the function of loading an operand into a register. In terms of their effect upon branching, they may also be called "mask" instructions, because they mask the conditioning effect of a previous CB instruction; they may perform functions other than loading a register. The mask emplaced by statement #15 is not automatically removed by the next instruction. Statement #16, for instance, cannot cause a branch for any result, since latch M is in a reset state during its execution. Statement #17, however, is a "stack logic" instruction and will therefore set latch M, since latch L is still set and SAR 272 contains all zeros. Statement #17 will then effect a branch (i.e., set latch B) if it produces an all-zero result in RR1 510. Masking and unmasking of a CB instruction may occur any number of times.

15 The primary purpose of masking is to ensure that a branch cannot be taken at a different parenthetical level of a logic equation than that which existed when the CB instruction occurred. In the present embodiment, such condition branching is restricted to the first level—i.e., when SAR 272 contains all zeros. If found to be desirable, it would obviously be possible to provide additional gating facilities for recording the stack level at which a CB instruction occurs, and then to allow that address to be taken only at the same stack level.

20 Another group of instructions, named "de-condition" instructions in Fig. 6, cause the effect of a previous CB instruction to be terminated absolutely. These instructions reset latches L and M. Although the previously specified branch address remains in BAR 225, it can never be transferred to PAR 223, since latch B can only be set when latch M is set, latch M can be set only when latch L is set, and only a subsequent CB instruction can set latch L. Such a subsequent CB instruction will, of course, contain its own branch address, which will overwrite any address previously

25 placed in BAR 225. De-condition instructions are generally those which themselves cause branching or linking, or which load specified registers.

30 Operand Indexing

35 As has been described in connection with Figs. 2 and 6, the operand addresses contained in some instructions are indexed before being applied to DAR 262 of DSM 260. Indexing is well known in conventional computer structures; in these conventional structures, however, address wraparound occurs at the end of the physical memory, and is usually unintentional. In the present system, an instruction within the program flow sets an artificial memory size as a modulus for address wraparound. Therefore, one or more selectable portions of DSM 260 appear to a program to be cyclically repeated.

40 Instructions having indexed operands, such as ACX, LX, OCX, etc., employ the INDEX gating function shown in Table I. This function expresses symbolically the operation of Index Unit 300 (as explained in connection with Fig. 3), followed by the transfer of SR 243 to DAR 262 via gate 263. The symbol "NB" in Fig. 7 is a collective name for the outputs 305—307 of decoder 657, Fig. 6. As has been explained, the value of NB is set by two control bits in an SXS instruction, which set latches C0 and C1 (collectively named "C" in Table I). In the present embodiment, the possible values for NB are "8", "9" and "10", which correspond to memory segments of 256, 512 and 1024 bytes, respectively. Other values could be chosen, and it would of course be possible to choose memory-segment lengths which are not integral powers of two, if desired. As explained in connection with register 650, Fig. 6, modular indexing may be used to confine the actual data operands fetched by indexed instructions to a particular area of memory. In the recognition of OCR A-Font characters, for example, the recognition logic need test only eighteen scans at any time. Therefore, operand addresses may be indexed from an arbitrary point within a 256-byte data area to any other point in the same area, without creating any programming overhead for ensuring that the absolute address produced by the indexing operation remains within the area.

45 The common function WHITEFILL, Fig. 7, describes the generation of the signal at the output of OR gate 657, Fig. 6, from the carry signal WC on line 308. As has been explained, the signal from OR 667 controls gate 26A into DOR 269. Turning now to the description in Fig. 7 of a typical indexed instruction, it will be seen that in the AX instruction, for instance, the primary logic function performed is an AND between RR1 510 and DOR 269. The latter quantity, however, is ANDed with WHITEFILL, so that the operand from DOR

269 is effectively made all zeros when there is not output from OR 667, regardless of the actual content of the DSM location addressed by DAR 262.

5 The combination of variable-modulus indexing and white filling allows processor 200 to focus on a particular portion of an input video pattern, while blanking out the remainder of the pattern. That is, this feature creates a 10 background-color mask having a variable-size window which may be placed over the pattern.

#### Voting Logic

15 The set of result registers 500, Fig. 5, provide a simple and rapid calculation of voting-logic ("m out of n", or "m/n") functions. Although circuit modules for performing these functions are known in the art, they have not found their way into the architecture of 20 general-purpose computers. The present invention, however, allows their inclusion at minimal additional cost.

25 Voting logic is performed by the programmer with the "Push" (P, PC, PCX, PX, PSC, PX) instructions, as described in Fig. 7. An instruction such as "L Z55", for instance, fetches the operand Z55 from DSM 260 and loads it in RR1 510. An instruction "P Z68" then places the "1/2" function of Z55 and 30 Z68 in RR1 510, and also places the "2/2" function of the same operands in RR2 520. A further instruction "P Z74" deposits the "1/3", "2/3" and "3/3" functions of all three 35 operands in RR1 510, RR2 520 and RR3 530, respectively. Additional instructions of the same type would place the "1/n", "2/n" and "3/n" functions of the n operands in the result registers. If desired, additional hardware result registers could be provided to 40 allow the calculation of the "4/n", "5/n", etc. voting functions. Three result registers allows the calculation of any voting function in which either m or n-m+1 is three or less, since the m/n function of variables X1, X2, . . . Xm is identical to the complement of the 45 (n-m+1)/n function of the variables X1, X2, . . . Xm. The "push complement" instructions perform the operand inversions with no time penalty.

50 After a sequence of push instructions, the voting-function results are available in the registers 510, 520 and 530 for direct use by other instructions. Normally, however, one or more of the results will be stored in DSM 260 or SSM 270 by a "Store" instruction such as 55 S, S2, S3, SS, S2S or S3S.

60 The usefulness of this implementation becomes even more apparent when it is considered that the "1/n" and the "n/n" voting functions are identical to the Boolean OR and AND functions, respectively. Thus, for instance, the OR and AND of three variables may be calculated and stored with only five instructions (one load, two push and two

store), whereas a conventional one-address processor would require at least eight instructions.

#### Data Input

65 Unbuffered video bytes from processing circuits 150 are transmitted to video register 281, and thence through gate 254 and DIR 253 into DSM 260. To feed such asynchronous data by cycle-steal or interrupt methods would tie up the entire processor memory for a data-entry operation and might also require additional programming steps. The use of dual parallel processors for this operation would 70 require a significant amount of additional hardware and complex interconnection and priority logic.

75 Instead of these more conventional techniques, processor 200 uses a quasi parallel-processor approach in which one of the processors is implemented as special-purpose hard-wired logic 630. This logic is responsive to a signal generated in instruction decoder 230 to perform in one instruction cycle a number of operations (each of which would require a full cycle in a fully programmed processor) for controlling gates 254, 261 and 263.

80 The hardware used in implementing this feature has been described hereinabove. The equivalent APL description is contained in the VIDEO function shown in Fig. 7. Those instructions in Fig. 7 which do not have the VIDEO function require operands from DSM 260; hence, no data input thereto is possible. The instructions which obtain their operands from other sources, however, generate a "video service" signal to logic 630, which initiates 85 the VIDEO function. Referring to Fig. 7, the AMZ instruction, for instance, causes a previously conditioned branch address to be effective upon the occurrence of a specified result of a logic operation. The operands for this instruction are obtained from RR2 520 and from POR 228. During the execution time of this instruction, logic 630 determines 90 whether VR 281 contains a new byte of video data, and whether DSM 260 contains a free location to store it. If both conditions exist, then the contents of VR 281 are gated into DSM 260, and the contents of TOP 264 are simultaneously transmitted to DAR 262. If 95 either of the conditions is not fulfilled, then 100 no data is entered into DSM 260.

105 110 115 The hardware used in implementing this feature has been described hereinabove. The equivalent APL description is contained in the VIDEO function shown in Fig. 7. Those instructions in Fig. 7 which do not have the VIDEO function require operands from DSM 260; hence, no data input thereto is possible. The instructions which obtain their operands from other sources, however, generate a "video service" signal to logic 630, which initiates the VIDEO function. Referring to Fig. 7, the AMZ instruction, for instance, causes a previously conditioned branch address to be effective upon the occurrence of a specified result of a logic operation. The operands for this instruction are obtained from RR2 520 and from POR 228. During the execution time of this instruction, logic 630 determines whether VR 281 contains a new byte of video data, and whether DSM 260 contains a free location to store it. If both conditions exist, then the contents of VR 281 are gated into DSM 260, and the contents of TOP 264 are simultaneously transmitted to DAR 262. If either of the conditions is not fulfilled, then no data is entered into DSM 260.

#### Automatic Result Stacking

120 The storage of intermediate results in a stack has been employed in conventional practice. The present invention, however, contains a simplified control mechanism which combines a conventional data memory with a hardware stack, and which stores data in the stack only when required to preserve the data. Moreover, selection between the stack and the random-access memory need not entail additional programmed instructions. This feature 125

is especially useful for the evaluation of multi-level logic equations.

When an explicit "store" instruction is employed in a program, the contents of RR1 510 are placed at a DSM address specified by the instruction. Therefore, RR1 510 may be overwritten by a subsequent instruction without totally destroying the previous result, since it has been preserved in DSM 260.

10 Hard-wired logic 610 performs this function by setting latch D ( $D \leftarrow 1$  in Fig. 7) whenever the contents of RR1 510 are stored. When no such storage has occurred, any "logic" instruction resets latch D ( $D \leftarrow 0$  in Fig. 7) so that a subsequent "load" instruction causes AND 611 to transfer the contents of RR1 510 through gates 254 and 271 to SSM 270. This operation is shown as the STACKIN function in Fig. 7. Since this transfer occurs before new data enter RR1 510, the previous contents are automatically saved. The "load" instructions also reset latch D after a delay. Thus the first "load" instruction following a store instruction overwrites RR1 510, but subsequent "load" instructions cause a transfer to SSM 270.

Data may also be stored in the stack under explicit program control. Instructions such as SLS, SS, S2S, etc. place the contents of RR1 510 in SSM 270 regardless of the state of latch D. These instructions also set latch D, since RR1 510 may safely be overwritten by subsequent instructions. It may be seen, then, that this feature provides both automatic and explicit control of data entry into SSM 270, and that automatic operation is provided without any software overhead.

#### Shift-Load Functions

The inclusion of shift gates 540 in result register unit 500 allows the performance of an unconventional function which considerably reduces processing time in certain applications. This function is executed by the "Shift Left Load Right" and "Shift Right Load Left" instructions LR(N) and LL(N) respectively, which also exist in the indexed-operand forms LRX(N) and LLX(N), and in the stack-operand forms LR(N)S and LL(N)S. The "N" in the mnemonic represents the number of bits to be shifted and loaded, and is encoded in two bit positions in the instruction word. In the present embodiment this number may be one, two or four.

The control-gating description for the above instructions is contained in Fig. 7. As a specific example of their operation, consider the instruction LR4. After the instruction is fetched, any previous CB instruction is de-conditioned by resetting latches L and M. The number (four in this case) of bits to be shifted and loaded is then obtained from bits four and five of the instruction in POR 228. Bits six through fifteen of POR 228 are transferred to SR 243 without indexing, and thence to

DAR which also receives the state of latch H in its high-order position. DSM 260 then outputs into DOR 269 the contents of the location specified by DAR 262, RR2 520 and RR3 530 are reset to zero. The contents of RR1 510 are end-around shifted by four places and transferred to RR2 520, so that positions (0, 1, 2, 3, 4, 5, 6, 7) of RR2 520 contain positions (4, 5, 6, 7, 0, 1, 2, 3) of RR1 510. The contents of bit positions (4, 5, 6, 7) of DOR 269 are then loaded into positions (4, 5, 6, 7) of RR1 510. Finally, PAR 223 is incremented by one by NEXT, since conditional branching has been disabled.

At the completion of any shift-load instruction, RR3 530 contains all zeros. RR2 520 contains the end-around shifted version of the byte previously stored in RR1 510, the amount and direction of the shift being controlled by certain bits of the instruction. RR1 510 retains a portion of its original contents, and also contains a portion of another operand. The instruction specifies the source of the operand and whether or not it is to be indexed. It will be noted at this point that a sequence of shift-load instructions may be employed to place bits from more than two operands into a single data byte in RR1 510.

#### Programming Examples

Table I shows a typical instruction sequence for recognizing a machine-printed character on a document. Each instruction in this source listing has a statement label, an operation code and an operand description. The operands beginning with M through Z represent certain "features" or subpatterns of the character, and may be derived by hardwired logic circuits, or by the evaluation of logic equations in processor 200 or in another special-purpose or general-purpose data processor. Operands beginning with A through L represent video bits taken directly from the input pattern stored in DSM 260.

The program segment in Table II evaluates the Boolean expression:

T0 = (PC . MAF2 . Z3 .  $\overline{Z64}$  .  $\overline{K9}$  .  $\overline{K10}$  . 110  
 $\overline{Z9}$  .  $\overline{Z80}$  .  $\overline{Z43}$  .  $\overline{Z11}$  .  $\overline{Z23}$  . Y311) .  
 $\overline{((YL11) . (Z68+Z55) . (Z45+Z21) . (Z58+Z74 . Y83 . Z13) . (Z49 . Z13+Z48)) . ((Z1+Z58 . Y21 . Z21 . Z42) . (YB45+Z118 . Z108)) . ((X246 . Z49+Z76) . (Z72+Z40) . 115  
 $\overline{(Z69+YR36)) . ((YO10+Z55) . (YB65))}$ ,  
which represents an OCR-A font numeral "2" in terms of the named features and video bits. That is, if the equation is satisfied, the character being classified may be a "2". Other 120$

equations for other characters must also be evaluated, until one and only one of them is satisfied.

The complexity and number of these equations points out the well-recognized deficiencies in the application of conventional computers to pattern recognition. The brute-force evaluation of all terms of each equation requires inordinate amounts of time, while the inclusion of testing instructions after the appropriate terms requires excessive amounts of storage. The overhead involved in data management and other auxiliary functions consumes even more time and/or storage.

15. The CB instruction at label #2 is extremely important in overcoming these problems. This single instruction is effective to cause a branch to statement #57 from any of the twenty-three statements #3—#14, #17, #20, #25, #29, #35, #39, #43, #46, #49, #52, and #53. As may be seen from the logic equation, these statements are points at which the value of the entire function can be tested. That is, the equation is false if any of the designated instructions returns a "zero" result, since they all involve logic AND operations. The possibility of branching is masked at statements #15, #18, #21, #26, #30, #36, #40, #44, #47 and #50 since the corresponding load instructions reset latch M and push an intermediate result into SSM 270. A subsequent instruction such as #16 cannot then cause a branch, since the stack remains pushed down until it is popped by the AS instruction at #17. Otherwise, the OR operation at #16 would cause a branch for a "one" results, since the truth of any OR'ed operand at the highest level of an equation establishes the value of the equation without further processing. The branch address established by the CB instruction is absolutely deconditioned by statement #54, which stores the value of the logic function. Thereafter, no conditional branch can be taken until another CB instruction sets up another address in BAR 225. If the CB instruction at statement #2 results in a branch, the execution of the branch also provides a deconditioning operation until another CB instruction occurs.

The program segment in Table I also illustrates the use of automatic stacking of intermediate results. Although the logic-equation operands PC, Z3, Z64, etc. are obtained in the conventional manner from the random-access DSM 260, Table II has no explicit commands for storing any result except the final value of the function. Instead, the load instructions at statements #15, #18, #21, #26, #30, #36, #40, #44, #47 and #50 automatically store the previous computational results in SSM 270. Statement #1, for instance resets latch D if it has been set by a previous store instruction, and loads operand PC into RR1 510. Statement #2

does not use RR1 510, so requires no storage of its contents. Statements #3—#14 perform successive AND operations, leaving an intermediate result in RR1 510. The load instruction at #15 transfers this result to SSM 270 before writing its operand in RR1 510, since latch D is reset. Latch D remains reset until statement #54 stores the contents of RR1 510 into DSM 260. A subsequent load command would then merely overwrite the previous contents of RR1 510.

Any instruction which stores data in SSM 270 also increments SAR 272, as is done in conventional practice. Likewise, the retrieval of data from SSM 270 decrements SAR 272. Data bytes are retrieved from the stack by "stack-logic" instructions such as statement #17. In actuality, operation-code space is saved by making the stack-logic operation codes the same as those for the corresponding ordinary logic instructions, the mnemonics are different only for ease of reading a program. The "stack-logic" instructions are differentiated from their counterpart logic commands in that the former have absolute DSM addresses of "0000", which is detected by decoder 231, Fig. 2.

Another function which occurs implicitly in the above program segment is the loading of video data from VR 281 into DSM 260. This operation may take place during the execution of statements #2, #17, #20, #25, #29, #35, #39, #43, #46, #49, #52 and #55—#57, since these instructions do not otherwise employ DSM 260. Thus, input data may be entered into processor 200 during about 25% of the instruction cycles shown. The data bytes are automatically fed into successive DSM locations by means of BOT 267, DAR 262 and incrementer 265. Registers TOP 264 and BOT 267 are initialized for each input character by an instruction of the form "ZTB K1, K2", where K1 and K2 are two four-bit immediate operands. As may be seen in Fig. 7, these operands set the four high-order bits of TOP 264 and BOT 267; the low order seven bits of these registers are reset to zero.

Statements #6 and #7 in Table I contain instructions having indexed operands. Suppose that a previous instruction "SXS 0, 0" had already set logic circuits 650 and 660 to select an eight-bit addition modulus and to disregard carries from IU 300. Also suppose that two other instructions "F BEGIN" and "LXR 16" had placed the sum of the symbolic variable "BEGIN" and the immediate data "16" into XR 244. Then the actual operand address used by statement #6 is the modulo-256 sum of the low-order eight bits of the address of K9, the value of the operand BEGIN and the constant "16", this sum then being added to 256 times the sum of the two high-order bits of the address of K9 and to 1024 times the value of the H bit.

Table II illustrates a program segment for creating shifted replicas of an input pattern. The video spots and measurements employed in pattern recognition normally have only two states: white or black, on or off, etc. Although a single-bit operand is capable of expressing these states, the data paths of execution unit 250 are eight bits wide. This feature allows the simultaneous examination of corresponding video bits or measurements from eight different patterns. In many applications of the present system it is desirable to execute the same recognition algorithm on each input pattern in several different locations on the pattern, to correct for size and positioning errors. The "shift-load" instructions described hereinabove offer a convenient and rapid method of generating patterns which are replicas of the original input pattern except that they are displaced one or more bit positions horizontally and/or vertically.

Table II shows a program sequence for generating an "8 by 1" expansion of a video pattern. That is, if the pattern is scanned with vertical raster scans having, e.g., fourteen bits per scan, then the  $8 \times 1$  expansion loads one horizontally adjacent bit from each of eight scans into a single byte in DSM 260. The operands in Table II are referenced to bits 1-8 of scan A, and the final results are located in the variables ZZ1-ZZ8. The ZZ1 byte originally contains the most recent eight bits from VR281, and is loaded into RR1 510 by an instruction preceding those shown in the table. It should be noted that the "store" and "load" instructions are of the indexed form SX and LX; thus, all scans of the pattern may be processed by this sequence. Again, an all-zero operand indicates that the operand is taken from SSM 270, rather than from DSM 260.

Other types of expansion are also possible with the "shift-load" instructions. A "4 by 2" expansion, for instance, uses the LL2 or LR2 instructions to enter two adjacent bits from each of four scans into a single byte. A "2 by 4" expansion loads four bits from each of two scans into each byte, using the LL4 or LR4 instructions. A "1 by 8" operation uses LL1 and LR1 instructions to place eight bits of a single scan into each byte.

Table III shows a simple application of processor 200 to the control of machine functions by means of the external registers 280. Although the above described embodiment of the processor has been optimized for the execution of pattern-recognition algorithms, it is also capable of acting as a machine controller. Machine control, in fact, is conceptually not completely remote from the field of pattern recognition, since both are frequently performed in real time, both involve a large number of data operands each having a relatively small number of bits, and both must frequently evaluate complex logic functions.

The instruction sequence in Table III reads EXR32 281 into RR1 510, and branches to a statement labelled OFF if either the first bit is a "zero" or the fourth bit is a "one". Otherwise, RR1 510 is reset to "zero" and written back into EXR32 282. Various bits of EXR32 282 may be connected to various sensors and control circuits of the system 100, Fig. 1, or in other systems.

This application together with United Kingdom Patent Application No. 22930/75 (Serial No. 1435928) was divided out of United Kingdom Patent Application No. 32161/73 (Serial No. 1435926), which applications claim different features of the apparatus disclosed herein.

TABLE I

85	#1	L	PC	#22	AC	Y83	105
	#2	CB	#57	#23	A	Z13	
	#3	A	MAF2	#24	O	Z58	
	#4	A	Z3	#25	AS	0000	
	#5	AC	Z64	#26	L	Z49	
90	#6	ACX	K59	#27	A	Z13	110
	#7	ACX	K9	#28	O	Z48	
	#8	AC	Z9	#29	AS	0000	
	#9	AC	Z80	#30	L	Z58	
	#10	AC	Z43	#31	A	Y21	
95	#11	AC	Z11	#32	A	Z21	115
	#12	AC	Z23	#33	A	Z42	
	#13	A	Y311	#34	O	Z1	
	#14	AC	YL11	#35	AS	0000	
	#15	LC	Z68	#36	LC	Z118	
100	#16	O	Z55	#37	AC	Z108	120
	#17	AS	0000	#38	OC	YB45	
	#18	L	Z45	#39	AS	0000	
	#19	O	Z21	#40	LC	X246	
	#20	AS	0000	#41	AC	Z49	
	#21	L	Z74	#42	OC	Z76	125

TABLE I (cont.)

5	#43	AS	0000	#51	O	Z55	10
	#44	L	Z72	#52	AS	0000	
	#45	O	Z40	#53	A	YB65	
	#46	AS	0000	#54	S	T0	
	#47	LC	Z69	#55	ZR	X'F2'	
	#48	O	YR36	#56	BAL	CTBL	15
	#49	AS	0000	#57	BAL	INIT	
	#50	L	Y010	#58	.....		

TABLE II

20	#1	SS	0000	#17	SX	ZZ3	35
	#2	SLS	0000	#18	LX	A4	
	#3	SLS	0000	#19	LL1S	0000	
	#4	SLS	0000	#20	SX	ZZ4	
	#5	SLS	0000	#21	LX	A5	
	#6	SLS	0000	#22	LL1S	0000	40
	#7	SLS	0000	#23	SX	ZZ5	
	#8	SLS	0000	#24	LX	A6	
	#9	LX	A1	#25	LL1S	0000	
	#10	LL1S	0000	#26	SX	ZZ6	
	#11	SX	ZZ1	#27	LX	A7	45
	#12	LX	A2	#28	LL1S	0000	
	#13	LL1S	0000	#29	SX	ZZ7	
	#14	SX	ZZ2	#30	LX	A8	
	#15	LX	A3	#31	LL1S	0000	
	#16	LL1S	0000	#32	SX	ZZ8	50

TABLE III

#1	READ	EXR32	#4	OMO	B'11101111'	55
#2	CB	OFF	#5	ZR	0	
#3	AMZ	B'10000000'	#6	WRT	EXR32	

## WHAT WE CLAIM IS:—

1. A data processor including apparatus for performing a plurality of different voting-logic operations such apparatus comprising: a plurality of cascaded registers; data storage means for holding and accessing a plurality of data operands; an instruction unit for accessing and interpreting program instructions, said instructions including a plurality of instructions respectively containing specifications of a corresponding plurality of said operands, and said program instructions including a further instruction containing a specification of a selectable one of said registers; means coupled to said instruction unit and to said data storage means for entering successive ones of said operands into a first of said cascaded registers and for transferring the contents of the register sequentially from register to register so that for any register its new contents are a logical combination, with respect to a particular fixed value, of its former contents and the former contents of the preceding register (or in the case of the first register, the new operand); and gating means coupled to said instruction unit and to all of said cascaded registers for transferring to a common output line the contents of said selectable one register, so as to gate out one of a plurality of different voting-logic functions of said successive operands, depending upon which of said cascaded registers is specified by said further instructions. 90
2. A processor as claimed in claim 1, wherein said particular fixed value is a logic "one".
3. A processor as claimed in claim 2, wherein said plurality of cascaded registers comprises at least three registers, thereby allowing said further instruction to selectively specify at least the voting-logic functions one out of n, two out of n and three out of n, where n is any integer. 95
4. A processor as claimed in claim 2, wherein each of said registers has a plurality of positions respectively adapted to hold single binary digits. 100
5. A processor as claimed in claim 4, wherein said instruction unit is further adapted to access a shift-load instruction containing a specification of a data operand; wherein said apparatus further comprises shifting means coupled to said instruction unit, to said data storage means and to two of said plurality of cascaded registers for shifting said binary digits from a particular one of said registers and for replacing predetermined ones of said

digits with a portion of said operand specified by said shift-load instruction.

5 6. A processor as claimed in claim 5, wherein said shifting means is coupled between said first register and to a second of said cascaded registers.

7. A data processor as claimed in claim 1,

substantially as hereinbefore described with reference to and as illustrated in the accompanying drawings.

10

I. M. GRANT,  
Chartered Patent Agent,  
Agent for the Applicant.

Printed for Her Majesty's Stationery Office, by the Courier Press, Leamington Spa, 1976.  
Published by The Patent Office, 25 Southampton Buildings, London, WC2A 1AY, from  
which copies may be obtained.

This drawing is a reproduction of  
the Original on a reduced scale  
Sheet 1




FIG. 1




FIG. 4




FIG. 2

1435927

## COMPLETE SPECIFICATION

14 SHEETS

This drawing is a reproduction of  
the Original on a reduced scale

Sheet 3





FIG. 5

FIG. 3



COMMON FUNCTIONS

THESE ROUTINES ARE COMMONLY USED WITHIN MANY INSTRUCTIONS.

```

▼ FETCH
  POR=PSN[21PAR1]
  OPR=POR[18]
  ▼

▼ Z+X PLUS Y
  Z+(10p2)T(21X)+21Y
  ▼

▼ NOINDEX
  SR=POR[6+110]
  DAR=H,SR
  ▼

▼ INDEX
  SRI[0 1]+POR[6 7]
  SRI[(10-NB)+1(NB)]+(NBp2)T(21POR[16-NB]+1NB)+21XR[(10-NB)+1NB]
  DAR=H,SR
  ▼

▼ DFETCH
  DOR=DSN[21DAR1]
  ▼

▼ DSTORE X
  DIR=X
  DSN[21DAR1]+DIR
  ▼

▼ NEXT
  +(B=1)/BRANCH
  SEQUENTIAL:PAR[4+112]+(12p2)T
  -1+21PAR[4+112]
  -END
  BRANCH:PAR+BAR
  L=N+0
  EBD:B=0
  ▼

▼ VIDEO
  +(V=0)VQ)/DONOTHING
  FEEDVIDEO:U=0
  DAR=BOT
  BOT[3+8]+(8p2)T(21BOT[3+16,1+1]
  Q=A/BOT=TOP
  DSTOR VR
  V=0
  DONOTHING:
  ▼

▼ PUSH X
  RR3=RR3+RR2AX
  RR2=RR2+RR1AX
  RR2=RR1VX
  ▼

▼ RESET
  RR3=RR2+8p0
  ▼

▼ STACKIN X
  +D=1/STACKEND
  DEQUALSO:SAR+(4p2)T1+21SAR
  DIR=X
  SSM[SAR1]+DIR
  STACKEND:D=0
  ▼

▼ Z+STACKOUT
  Z+SSM[SAR1]
  SAR+(4p2)T(21SAR)-1
  ▼

▼ INITIALIZATION
  D+U+V+B+L+N=0
  H=0
  C=0 0
  NB=8
  ▼

▼ Z+WHITEFILL
  WC=(1,8Bp0)/((NB+1)p2)T(21IR[10-NB]+1NB)+21POR[(16-NB)+1NB]
  Z+8p~(W[0]A=WC)VW[1]AWC
  ▼

```

FIG. 7

INSTRUCTIONS

SEE COMMON FUNCTIONS FOR ROUTINES USED WITHIN INSTRUCTIONS.

ADD

```

V A
  FETCH
  NOINDEX
  DPETCH
  RR2~RR1
  RR1+RR2ADOR
  B~NA~V/RR2ADOR
  RESET
  D=0
  NEXT
  V

```

ADD

```

V ADD
  FETCH
  SR+RR2 PLUS (POR [2+18] PLUS POR [12])
  RR1~SR[2+18]
  B~NA~POR[15]~SR[1]
  Z~(0..1,7p0)/(RR2[1+17] PLUS
  (POR [3+17] PLUS POR [12]))
  B~B~NA~POR[14]~Z
  B~B~NA~POR[13]~SR[1]#2
  VIDEO
  NEXT
  V

```

ARITHMETIC ADD

```

V AA
  FETCH
  SR+POR[8+18] PLUS POR[2+18]
  B~NA~SR[1]
  RR1~SR[2+18]
  VIDEO
  NEXT
  V

```

ADD MASK NOT ZERO

```

V ANZ
  FETCH
  RR2~RR1
  B~NA~V/RR2~POR[8+18]
  RR1+RR2~POR[8+18]~BpN
  VIDEO
  NEXT
  V

```

AND COMPLEMENT

```

V AC
  FETCH
  NOINDEX
  DPETCH
  RR2~RR1
  RR1+RR2~DOR
  B~NA~V/RR2~DOR
  RESET
  D=0
  NEXT
  V

```

AND MASK ZERO

```

V ANZ
  FETCH
  RR2~RR1
  B~NA~V/RR2~POR[8+18]
  RR1+RR2~POR[8+18]~BpN
  VIDEO
  NEXT
  V

```

AND COMPLEMENT INDEXED

```

V ACIX
  FETCH
  INDEX
  DPETCH
  RR2~RR1
  RR1+RR2~DOR~WHITEFILL
  B~NA~V/RR2~DOR~WHITEFILL
  RESET
  D=0
  NEXT
  V

```

AND STACK

```

V AS
  FETCH
  RR2~RR1
  RR1+RR2~STACKOUT
  N~N~LA0~V/SAR
  B~NA~V/RR1
  VIDEO
  RESET
  D=0
  NEXT
  V

```

ABD STACK COMPLEMENTED

▼ ASC  
FETCH  
RR2→RR1  
RR1→RR2 $\wedge$ ~STACKOUT  
M→N $\wedge$ LA0 $\wedge$ v/SAR  
B→N $\wedge$ v/RR1  
VIDEO  
RESET  
D→0  
NEXT

▼

ABD INDEXED

▼ AX  
FETCH  
INDEX  
DFETCH  
RR2→RR1  
RR1→RR1 $\wedge$ DOR $\wedge$ WHITEFILL  
B→N $\wedge$ v/RR2 $\wedge$ DOR $\wedge$ WHITEFILL  
RESET  
D→0  
NEXT

▼

BRANCH

▼ B  
FETCH  
B→1  
L→M→0  
BAR[4+112]+POR[4+112]  
VIDEO  
NEXT

▼

BRANCH AND LINK

▼ BAL  
FETCH  
B→1  
L→M→0  
BAR[4+112]+POR[4+112]  
BLR $\leftarrow$ (16 $\times$ 2) $\wedge$ 1 $\wedge$ PAR  
VIDEO  
NEXT

▼

CONDITION BRANCH

▼ CB  
FETCH  
L→M→1  
BAR[4+112]+POR[4+112]  
VIDEO  
SAR $\leftarrow$ 4 $\times$ 0  
NEXT

▼

COMPARE MASK EQUAL

▼ CNE  
FETCH  
RR2→RR1  
B→N $\wedge$ v/RR2 $\wedge$ POR[8+18]  
RR1 $\leftarrow$ ((8p $\sim$ N) $\wedge$ RR2 $\wedge$ POR[8+18]) $\vee$ (RR2 $\wedge$ 80H)  
VIDEO  
NEXT

▼

COMPARE MASK NOT EQUAL

▼ CMNE  
FETCH  
RR2→RR1  
B→N $\wedge$ v/RR2 $\wedge$ POR[8+18]  
RR1 $\leftarrow$ ((8p $\sim$ N) $\wedge$ RR2 $\wedge$ POR[8+18]) $\vee$ (RR2 $\wedge$ 80H)  
VIDEO  
NEXT

▼

FETCH

▼ F  
FETCH  
NOINDEX  
DFETCH  
POR[2+18] $\wedge$ DOR  
RR2→RR1  
NEXT

▼

FETCH HIGH

▼ FH  
NOINDEX  
DFETCH  
POR[0,1] $\wedge$ DOR[6,7]  
RR2→RR1  
NEXT

▼

HALT

▼ HALT  
FETCH

▼

LOAD

▼ L  
FETCH  
M→0  
NOINDEX  
DFETCH  
RESET  
STACK $\leftarrow$  RR1  
RR1 $\leftarrow$ DOR  
NEXT

▼

LOAD BRANCH ADDRESS LOW AND BRANCH

✓ LBB  
FETCH  
B+1  
L+N=0  
NOINDEX  
DPETCH  
BAR[8+18]-POR[2+18]-DOR  
NEXT  
▼

LINK CONDITION

✓ LKC  
FETCH  
L+N=0  
PAR+BLR  
VIDEO  
+(0=v/RR1)/v  
NEXT  
▼

LOAD BRANCH ADDRESS HIGH

✓ LBB  
FETCH  
NOINDEX  
DPETCH  
BAR[18]-POR[2+18]-DOR  
NEXT  
▼

SHIFT RIGHT LOAD LEFT INDEXED (N)

✓ LLX(N)  
FETCH  
L+N=0  
H+2\*(POR[5]+2\*POR[4])  
INDEX  
DPETCH  
RESET  
RR2+RR1[(8-H)+1N],RR1[18-H]  
RR1+DOR[1N],RR2[N+18-H]  
NEXT  
▼

LOAD COMPLEMENT

✓ LC  
FETCH  
N=0  
NOINDEX  
DPETCH  
RESET  
STACKIN RR1  
RR1~~DOR  
NEXT  
▼

SHIFT RIGHT LOAD LEFT (N)

✓ LL(N)  
FETCH  
L+N=0  
H+2\*(POR[5]+2\*POR[4])  
NOINDEX  
DPETCH  
RESET  
RR2+RR1[(8-H)+1N],RR1[18-H]  
RR1+DOR[1N],RR2[N+18-H]  
NEXT  
▼

LOAD COMPLEMENT INDEXED

✓ LCX  
FETCH  
H=0  
INDEX  
DPETCH  
RESET  
STACKIN RR1  
RR1~~DOR~WHITEPILL  
NEXT  
▼

SHIFT RIGHT LOAD LEFT (N) FROM STACK

✓ LL(N)S  
FETCH  
L+N=0  
H+2\*(POR[5]+2\*POR[4])  
RESET  
RR2+RR1[(8-H)+1N],RR1[18-H]  
RR1~STACKOUT[1N],RR2[N+18-H]  
VIDEO  
NEXT  
▼

LINK

✓ LK  
FETCH  
L+N=0  
PAR+BLR  
VIDEO  
▼



MOVE INDEX REGISTER

✓ **MIR**  
 FETCH  
 SR+POR[8+18] PLUS XR  
 RR1+SR[2+18]  
 $B \cdot M \wedge (0,1,8p0) / POR[8+18]$  PLUS XR  
 $XR[10-NB] + 1 \cdot MB] \cdot SR[(10-NB) + 1 \cdot MB]$   
 VIDEO  
 NEXT  
 ✓

NO OPERATION

✓ **NOOP**  
 FETCH  
 VIDEO  
 NEXT  
 ✓

OR

✓ **O**  
 FETCH  
 NOINDEX  
 DFETCH  
 RR2+RR1  
 $RR1 \cdot RR2 \vee DOR$   
 $B \cdot MAA / RR2 \vee DOR$   
 RESET  
 $D \cdot 0$   
 NEXT  
 ✓

OR COMPLEMENT

✓ **OC**  
 FETCH  
 NOINDEX  
 DFETCH  
 RR2+RR1  
 $RR1 \cdot RR2 \vee \sim DOR$   
 $B \cdot MAA / RR2 \vee \sim DOR$   
 RESET  
 $D \cdot 0$   
 NEXT  
 ✓

OR COMPLEMENT INDEXED

✓ **OCX**  
 FETCH  
 INDEX  
 DFETCH  
 RR2+RR1  
 $RR1 \cdot RR2 \vee \sim DOR \cdot WHITEFILL$   
 $B \cdot MAA / RR2 \vee \sim DOR \cdot WHITEFILL$   
 RESET  
 $D \cdot 0$   
 NEXT  
 ✓

OR MASK NOT ONES

✓ **ONNO**  
 FETCH  
 RR2+RR1  
 $B \cdot M \wedge \sim A / RR2 \vee POR[8+18]$   
 $RR1 \cdot RR2 \vee POR[8+18] \wedge 8p-N$   
 VIDEO  
 NEXT  
 ✓

OR MASK ONES

✓ **ONO**  
 FETCH  
 RR2+RR1  
 $B \cdot MAA / RR2 \vee POR[8+18]$   
 $RR1 \cdot RR2 \vee POR[8+18] \wedge 8p-N$   
 VIDEO  
 NEXT  
 ✓

OR STACK

✓ **OS**  
 FETCH  
 RR2+RR1  
 $RR1 \cdot RR2 \vee \sim STACKOUT$   
 $M \cdot MVLAD = \sim / SAR$   
 $B \cdot MAA / RR1$   
 VIDEO  
 RESET  
 $D \cdot 0$   
 NEXT  
 ✓

OR STACK COMPLEMENTED

✓ **OSC**  
 FETCH  
 RR2+RR1  
 $RR1 \cdot RR2 \vee \sim \sim STACKOUT$   
 $M \cdot MVLAD = \sim / SAR$   
 $B \cdot MAA / RR1$   
 VIDEO  
 RESET  
 $D \cdot 0$   
 NEXT  
 ✓

OR INDEXED

✓ **OI**  
 FETCH  
 INDEX  
 DFETCH  
 RR2+RR1  
 $RR1 \cdot RR2 \vee DOR \cdot WHITEFILL$   
 $B \cdot MAA / RR2 \vee DOR \cdot WHITEFILL$   
 RESET  
 $D \cdot 0$   
 NEXT  
 ✓

PUSH

✓ P  
FETCH  
M=0  
NOINDEX  
DFETCH  
PUSH DOR  
D=0  
NEXT  
✓

PUSH COMPLEMENT

✓ PC  
FETCH  
M=0  
NOINDEX  
DFETCH  
PUSH ~DOR  
D=0  
NEXT  
✓

PUSH COMPLEMENT INDEXED

✓ PCX  
FETCH  
M=0  
INDEX  
DFETCH  
PUSH ~DOR~WHITEFILL  
D=0  
NEXT  
✓

PUSH FROM STACK

✓ PS  
FETCH  
M=0  
PUSH STACKOUT  
VIDEO  
D=0  
NEXT  
✓

PUSH FROM STACK COMPLEMENTED

✓ PSC  
FETCH  
M=0  
PUSH ~STACKOUT  
VIDEO  
D=0  
NEXT  
✓

PUSH INDEXED

✓ PI  
FETCH  
M=0  
INDEX  
DFETCH  
PUSH DOR~WHITEFILL  
D=0  
NEXT  
✓

EQUALS

✓ Q  
FETCH  
NOINDEX  
DFETCH  
RR2~RR1  
RR1~RR2~DOR  
B~MAA/RR2~DOR  
RESET  
D=0  
NEXT  
✓

EQUALS FROM STACK

✓ QS  
FETCH  
RR2~RR1  
RR1~RR2~STACKOUT  
M=M<sub>1</sub>L<sub>AO</sub>=V/SAR  
B~MAA/RR1  
VIDEO  
RESET  
D=0  
NEXT  
✓

EQUALS INDEXED

✓ QI  
FETCH  
INDEX  
DFETCH  
RR2~RR1  
RR1~RR2~DOR~WHITEFILL  
B~MAA/RR2~DOR~WHITEFILL  
RESET  
D=0  
NEXT  
✓

READ

▼ READ  
FETCH  
L+M+0  
RR1←EX[21POR[8+18];]  
RESET  
VIDEO  
NEXT  
▼

STORE RESULT

▼ S  
FETCH  
L+M+0  
NOINDEX  
DSTORE RR1  
D+1  
NEXT  
▼

SET ADDRESSING STATUS

▼ SAS  
FETCH  
H+POR[15]  
VIDEO  
NEXT  
▼

SHIFT LEFT STORE STACK

▼ SLS  
FETCH  
L+M+0  
RR1←RR1[1+17],RR1[0]  
D+0  
STACKIN RR1  
D+1  
NEXT  
▼

SHIFT RIGHT STORE STACK

▼ SRS  
FETCH  
L+M+0  
RR1←RR1[7],RR1[17]  
D+0  
STACKIN RR1  
D+1  
NEXT  
▼

STORE RESULT IN STACK

▼ SS  
FETCH  
M+0  
D+0  
STACKIN RR1  
D+1  
NEXT  
▼

STORE RESULT INDEXED

▼ SX  
FETCH  
L+M+0  
INDEX  
DSTORE RR1  
D+1  
NEXT  
▼

SET INDEXING STATUS

▼ SXS  
FETCH  
L+M+0  
C+POR[14,15]  
NB+8+(21C)  
H+POR[10,11]  
VIDEO  
NEXT  
▼

STORE MAJORITY 2

▼ S2  
FETCH  
L+M+0  
NOINDEX  
D+1  
DSTORE RR2  
NEXT  
▼

STORE MAJORITY 2 IN STACK

▼ S2S  
FETCH  
M+0  
D+0  
STACKIN RR2  
D+1  
NEXT  
▼

STORE MAJORITY 2 INDEXED

▼ S2X  
FETCH  
L+M=0  
INDEX  
DSTORE RR2  
D+1  
NEXT  
▼

STORE MAJORITY 3

▼ S3  
FETCH  
L+M=0  
NOINDEX  
DSTORE RR3  
D+1  
NEXT  
▼

STORE MAJORITY 3 IN STACK

▼ S3S  
FETCH  
M=0  
D=0  
STACKIN RR3  
D+1  
NEXT  
▼

TEST ONE INDEXED

▼ TOX  
FETCH  
INDEX  
DFETCH  
B+MA~(1,7p0)/DORWHITEFILL  
NEXT  
▼

TEST ONE INDEXED. LINK

▼ TOXL  
FETCH  
INDEX  
DFETCH  
B+MAA~(1,7p0)/DORWHITEFILL  
+(A=1)/NEXT  
L+M=0  
PAR+BLR  
+V  
NEXT: NEXT  
▼

TEST ZERO INDEXED

▼ TZX  
FETCH  
INDEX  
DFETCH  
B+MA~(1,7p0)/DORWHITEFILL  
NEXT  
▼

TEST ZERO INDEXED. LINK

▼ TZXL  
FETCH  
INDEX  
DFETCH  
B+MAA~(1,7p0)/DORWHITEFILL  
+(A=1)/NEXT  
L+M=0  
PAR+BLR  
+V  
NEXT: NEXT  
▼

WRITE

▼ WRT  
FETCH  
L+M=0  
EX[21POR[8+18];]~RR1  
VIDEO  
NEXT  
▼

EXCLUSIVE OR

▼ XOR  
FETCH  
NOINDEX  
DFETCH  
RR2~RR1  
RR1~RR2~DOR  
B+MA~v/RR2~DOR  
RESET  
D=0  
NEXT  
▼

EXCLUSIVE OR FROM STACK

▼ XORS  
FETCH  
RR2~RR1  
RR1~RR2~STACKOUT  
M+MVLAA=v/SAR  
B+MA~v/RR1  
VIDEO  
RESET  
D=0  
NEXT  
▼

EXCLUSIVE OR INDEXED

✓ **XORX**  
**FETCH**  
**INDEX**  
**DFETCH**  
**RR2=RR1**  
**RR1+RR2=DORAWHITEFILL**  
 $B \leftarrow M \leftarrow V / BR2 \neq DORAWHITEFILL$   
**RESET**  
**D=0**  
**NEXT**

INITIALIZE BRANCH ADDRESS HIGH

✓ **ZBH**  
**FETCH**  
 $BAR[18] \leftarrow POR[8+18]$   
**VIDEO**  
**NEXT**

INITIALIZE RESULT REGISTER ONE

✓ **ZR**  
**FETCH**  
 $M \leftarrow L \leftarrow 0$   
**RESET**  
 $RR1 \leftarrow POR[8+18]$   
**D=0**  
**VIDEO**  
**NEXT**

INITIALIZE TOP AND BOTTOM

✓ **ZTB**  
**FETCH**  
 $L \leftarrow M \leftarrow 0$   
 $TOP \leftarrow POR[8+14], 7p0$   
 $BOT \leftarrow POR[12+14], 7p0$   
 $Q \leftarrow V \leftarrow 0$   
 $U \leftarrow A / TOP = BOT$   
**NEXT**

INITIALIZE INDEX REGISTER

✓ **ZIR**  
**FETCH**  
 $L \leftarrow M \leftarrow 0$   
 $XR \leftarrow 0, 0, POR[8+18]$   
**VIDEO**  
**NEXT**